Reply on RC2

William Colgan et al.

Author comment on "Greenland Geothermal Heat Flow Database and Map (Version 1)" by William Colgan et al., Earth Syst. Sci. Data Discuss., https://doi.org/10.5194/essd-2021-290-AC3, 2022

---

reviewer comments in italic │ **author comments in bold**

This paper presents a database for heat flow measurements in Greenland. The database adds some more points and additional context to previous databases, so I think that the presentation of the database is a very useful contribution.

The paper then presents an analysis and discussion of the data, and presents a map of heatflow (Fig. 3) that is based on the data. Here, the authors use quite a bit of discretion and judgment in order to perform the interpolation, and this is necessary because the Greenland heat flow data are so unevenly distributed – almost no data in the interior of Greenland, and then sparse and variable data around the periphery. The authors do an excellent job of comparing their map to the heat flow maps from previous studies (e.g., Fig. 13), and they provide an excellent and informative discussion (Section 4) of the issues surrounding the collection, interpretation, and interpolation of the Greenland heat flow data.

Overall, I think this is an interesting and useful contribution that could be published after some revision. The paper is very well written and the figures are well-crafted and convincing. My one major concern is the authors' treatment of the NGRIP station, which I discuss below and I think it should be addressed upon revision. I also have several other specific points that I discuss below, and I think that addressing these will improve the impact of the work.

We thank the reviewer for their generally positive response to this work. We address the NGRIP anomaly in detail below.

Regarding NGRIP:

This station provides one of the only datapoints in central Greenland. It is thus quite valuable, but it is also problematic because it shows a large heat flow value (~130 mW/m^2) that is somewhat larger than other values. The authors here chose to disregard this point from their machine learning analysis (Fig. 3) although they did include this point in some tests with their jackknifing approach (Fig. 9). It seems to me that including or not including NGRIP station has dramatically large impacts on the resulting heat flow map (as can be seen by comparing the different maps in Fig. 13). This seems to indicate a larger uncertainty about the heat flow in central Greenland than the authors have expressed in their analysis – I think that they are underestimating the uncertainty about heat flow in central Greenland. For example, if they had included the possibility of high heat flow at

NGRIP station (maybe NGRIP values with large uncertainty associated with them), then they would have a significantly higher max GHF values in Fig. 3c.

The reviewer is correct that the inclusion of NGRIP has a disproportionately large impact on the machine learning algorithm. Clearly, the algorithm identifies a large area of Central North Greenland that is geologically similar to NGRIP, based on the twelve input geophysical datasets. When NGRIP is excluded, this area is assigned lower heat flow characteristic of other subglacial sites. When NGRIP is included, however, this area is assigned the higher heat flow of NGRIP. This leads to our assumption that the measured heat flow at NGRIP likely results from a local phenomenon that is not captured in our twelve input geophysical datasets. We now describe this more explicitly in Section 3.2. We now also include both the "with NGRIP" and "without NGRIP" data products available online (https://doi.org/10.22008/FK2/F9P03L), and we also present the "with NGRIP" simulation in similar detail to the existing "without NGRIP" simulation (Figure 3). This allows the user to decide whether to include or exclude the NGRIP anomaly.

I suggest that the authors find a way to incorporate the additional uncertainty about heat flow in central Greenland that is expressed by their own uncertainty about what the NGRIP value actually represents. There might be several ways to do this. In principle, the machine learning algorithm could be trained to be smart enough to recognize an out-of-range station and disregard it to some extent if necessary. Presumably there are stations elsewhere in the world that are similarly spurious but included in the analysis, and the machine learning algorithm can learn how to deal with them. But this might be too much for a revision of the current study. Instead, the authors could run the machine learning algorithm again but including NGRIP and take some sort of average of the estimates with and without NGRIP. Alternatively, they might present their jackknifing analysis with NGRIP in a bit more detail, to get an estimate of how large the uncertainty associated with central Greenland really is. (see my comment for Figure 9 below)

We now highlight the influence of the NGRIP anomaly on heat flow uncertainty by comparing the jackknifing ensemble spreads of the "with NGRIP" and "without NGRIP" in a new Figure 9. This clearly shows that the inclusion of NGRIP introduces a region of elevated heat flow uncertainty around NGRIP. The manuscript also notes that in a pre-processing step we exclude all global heat flow measurements > 200 mW/m2, as these are likely caused by local phenomenon. We cannot resolve the precise nature of the NGRIP anomaly, we can simply run simulations with and without the NGRIP data point and speculate that it is likely associated with local processes.

Overall, we really do not know much about the heat flow in central Greenland, and so it seems strange to throw away the one data point that we have from this region. Instead, it would be better to incorporate this uncertainty over the NGRIP point into larger uncertainty estimate for central Greenland.

We now provide the "with NGRIP" data product and associated uncertainty estimate.

Specific points related to NGRIP:
Line 323 – Here the authors present an argument for excluding the NGRIP station (with its very high heat flow) from the machine learning training data. It is true that the heat flow estimates in central Greenland are very sensitive to the observation at NGRIP – this is because this point is much isolated from the others. I do not think that this makes a good argument for excluding the point – instead data points from sparsely-covered areas would be *more* valuable and important to include. I think the authors should develop some sort of general rule for excluding or including points in their analysis, for example based on proximity to other points that could indicate if a given observation is representative of its region. Otherwise, it seems like they are picking and choosing which points to include

(and see my point about the next line).

We now clarify that the spatial location of a measurement is not as important in machine learning as in a conventional spatial interpolation. From a machine learning perspective, it is the relations between data points and input fields that is of most importance. This means that a data point does not become valuable simply due to location. In this sense, NGRIP is not a spatial outlier, but rather a geophysical outlier. Simply put, heat flow at NGRIP is not consistent with other heat flow observations in similar geological settings. In terms of a general rule for exclusion, we only adopt the exclusion of heat flow measurements > 200 mW/m2 from Lösing and Ebbing (2021). While we recommend the exclusion of NGRIP, we also provide a data product version that includes NGRIP.

Line 324 – The authors suggest that this point might not be statistically representative of the broader region – the authors have no way of knowing this, because there are no other measurements from this region. By this principle, other points that stand by themselves should be similarly dismissed from the database. This would remove pretty much all the points under the main ice sheet (Fig. 1), and there would be very little data left from interior Greenland. Instead, the main reason the authors are disregarding NGRIP is because of its high value – if it had been more "normal" then they would have included it. This is a bit dangerous territory, since excluding data points because they seem spurious can get subjective – and indeed this decision has a huge influence over the resulting heat flow map.

We now show that the influence of the NGRIP anomaly on heat flow uncertainty is disproportionately large by comparing the jackknifing ensemble spreads of the "with NGRIP" and "without NGRIP" in a new Figure 9. This clearly highlights that of all the on-shore measurements, the machine learning algorithm is most sensitive to NGRIP. While we continue to argue that these ensemble spreads justify regarding NGRIP as an outlier, we also now include full analysis and data product for the "with" NGRIP simulation.

Figure 4 – here the relative importance of the different input variables used for the machine learning algorithm are presented. But the most important one for continental Greenland is omitted – the decision to omit the NGRIP station. It seems like this decision should somehow be expressed in a figure like this.

We note that this figure reflects the internal structure of the trained machine learning and does not include pre-processing, such as decisions of which geophysical inputs or training data to include. We now note that the inclusion or exclusion of NGRIP does not fundamentally shift this importance ranking of input geophysical datasets.

Line 397 – Here the authors suggest that some of the extra heat flow at NGRIP may be due to hydrological processes. But wouldn't the associated hydrological processes likely indicate high overall heat flow from this region? That seems to be the case with the other hot springs discussed in the paper (Fig. 6a), and my understanding of hot springs in general (they usually are located in high heat flow areas). (see also my comment below for line 493)

We now clarify that local hydrological processes such as subglacial water flow or hot springs are sub-grid cell processes relative to the global scale of our machine learning algorithm. But here, we also caution that the available heat flow measurements generally do not support high regional heat flow where hot springs are found. For example, the hot springs along 70N in both East and West Greenland are not associated with high local heat flow measurements.

Other specific points within the paper:
Line 40 – The authors present a list of reasons that good heat flow information is necessary. I agree with this list, but I would also add that heat flow data provides useful

constraints on the thermal structure of the lithosphere: its elastic thickness, density of heat producing elements in the crust, etc. I think it would be useful to add this to the list, so as to also make the paper relevant for tectonophysicists.

We have now included that improved geothermal knowledge help constrain the thermal structure and properties of the lithosphere.

Table 2 – I think that all of the fields specified in the database are useful. There are uncertainties specified for all the components, except for the parameters that go into computing heat flow – namely the temperature gradient and the conductivity. I would think this information could be useful to those using the dataset. For example, if a user feels that the uncertainty in conductivity should be higher (e.g., if they have measurements that suggest this) then they could develop their own uncertainty measure.

Unfortunately, uncertainties in gradient and conductivity are very seldom reported. Indeed, there are many sites for which gradient and conductivity themselves are not reported, it is simply heat flow that is reported. We do, however, report when conductivity is assumed, rather than measured, in the comment section of a site, which provides a qualitative flag. It does not presently appear possible to systematically assess empirical uncertainties in conductivity and gradient at site level.

Table 2 - I also do not fully understand what is meant by "where only gradient or conductivity is reported" (in the statement about heat flow uncertainty). How would heat flow be computed if only one of these is reported? Do the authors mean "if only uncertainty in the temperature gradient or the conductivity is reported"? But in that case, why assign the uncertainty to a set value (e.g., 10%) and not simply use the reported values?

We have now clarified that this ±10% assumed uncertainty relates to sites where temperature gradient is reported, but thermal conductivity is assumed. This contrasts with ±5% assumed uncertainty at sites where both the gradient and conductivity are reported (i.e. no assumptions). While we use reported uncertainties where available, the vast majority of sites have no reported uncertainty, which compels our fractional uncertainty system. This uncertainty system is clearly imperfect, but it is sufficiently transparent and systematic to allow subsequent users to modify uncertainty assumptions.

Line 127 – I do not understand what is meant by "diminishing extreme values from surveys conducted in the late 1970s and early 1980s". Which are the extreme values – heat flow, conductivity, or temperature gradient? It seems that the authors are identifying these problematic points as having abs((dT/dz)*k – q) > 2, so why do they need to provide another explanation as "diminishing extreme values" – which doesn't seem to have a real meaning (in my understanding). Incidentally, in these instances it seems that dT/dz, k, or q could be reported wrongly, and the authors here are assuming that it is q that is the bad value. Isn't it equally likely that it is k or dT/dz? It seems that the authors should at least consider this possibility if they are replacing a value of q that was reported from a previous study. In any case, I think a bit more explanation here would help.

We now clarify that the majority of these reassessments are "down revising" extreme values from the 1970s and 1980s. This sentence is not meant to explain a selection criteria, but rather summarize the cohort of reassessed values. We feel it is important to explain that these reassessments generally pertain to older measurements that yielded high heat flow values along the mid-Atlantic ridge. Simply put, there is a clear spatial and temporal coherence to the lower quality IHFC data that we reassess. We also clarify that k and dT/dz are the primary measurements and that q is a secondary derived product.

Line 133 – I would say "lower resolution" instead of "relatively low resolution" since resolution is a relative quantity.

We have made this correction and now provide the respective horizontal resolutions of the BedMachine and ETOPO1 DEMs for quantitative comparison.